Release Date: September, 2015

Updates:

Database Foundations

1-4
Relational Databases and Normalization

ORACLE ACADEMY

Roadmap

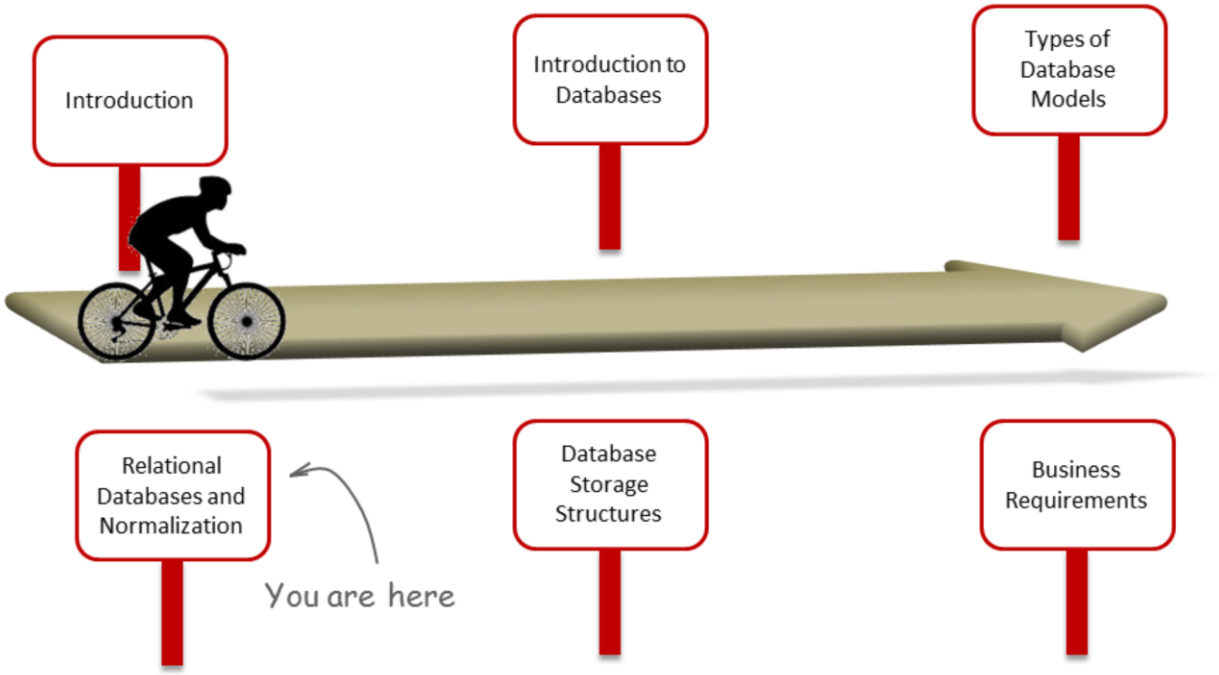

Objectives

This lesson covers the following objectives:

- Describe the features of a relational database
- Explain the rules of a relational database
- Explain the objectives of normalization
- Describe the types of normalization

Introduction to Relational Databases

- A relational database stores information in tables with rows and columns.
- A table is a collection of records.
- A row is called a record (or instance).
- A column is referred to as a field (or attribute).

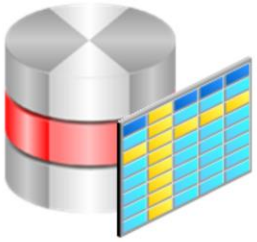

A relational database is a collection of records that are stored in tables. Each relational database table contains rows of records and columns containing fields of information about each record. A table in a relational database can also be referred to as an entity. A row in a relational database can also be referred to as an instance.

Each table of records will have a relationship to another table of records when the two tables share a field (or column).

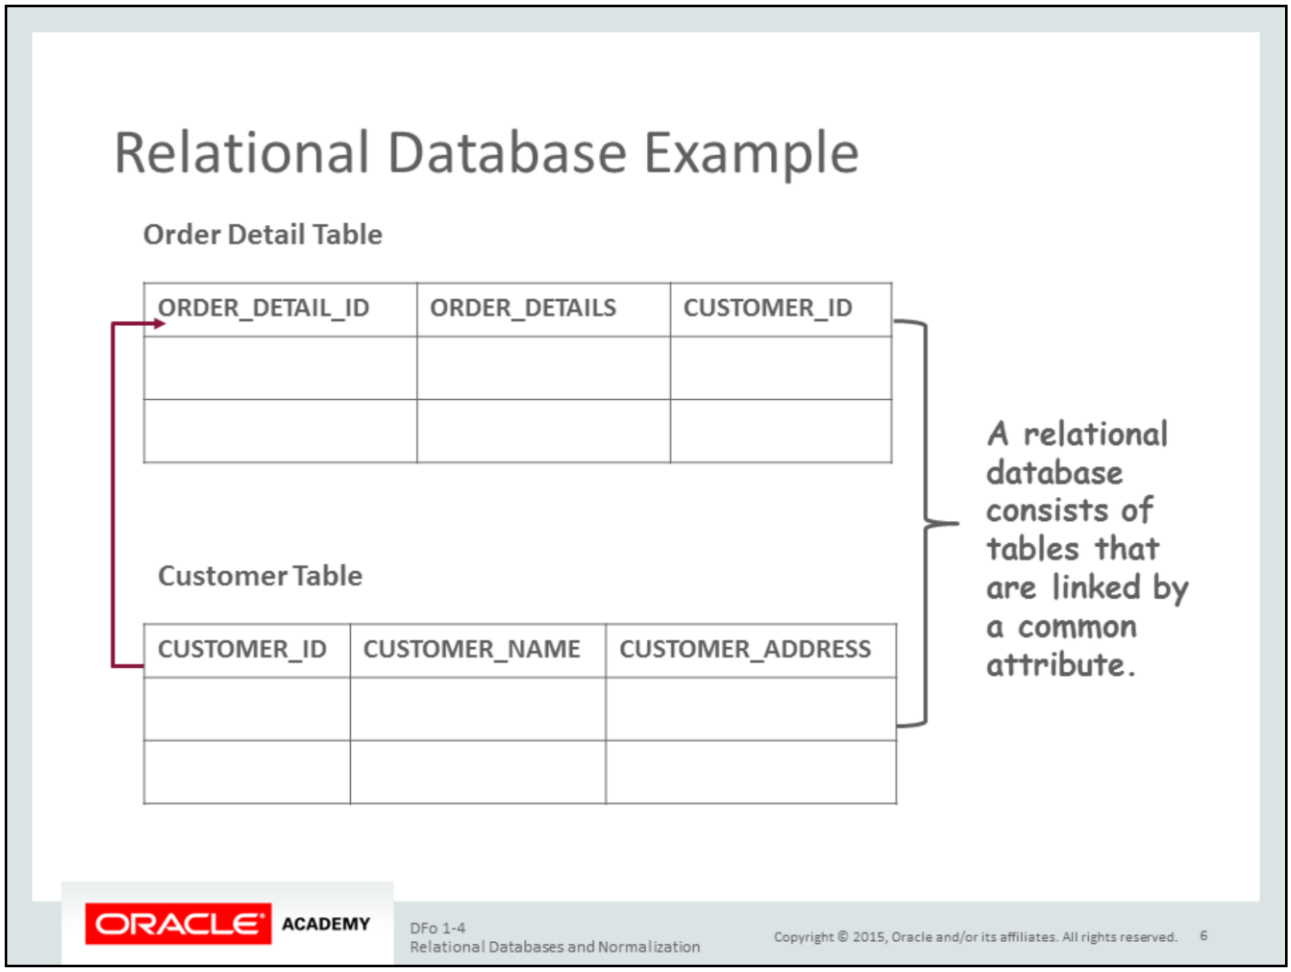

The slide depicts two tables: Order Details and Customer. The tables are related to each other by a common attribute, Customer ID.

Imagine a single order placed by a customer. Each order will contain one or more order details. Each detail will be related to one customer.

The data provides information about the details of orders placed by customers. For example, the company could gather information about products that are commonly purchased together. Bundles of products could then be offered to better market the products to customers.

Rules for Relational Databases

- Each table has a distinct name.
- Each table may contain multiple rows.
- Each table has a value to uniquely identify the rows.
- Each column in a table has a unique name.

The basic rules listed in the slide are implemented as part of the overall relational database design.

When you design a relational database, it is important that you minimize redundant data (that is, the same data stored in more than one place in the database.)

Normalization

- Is the process of organizing the attributes and tables of a relational database to minimize redundancy.
- Helps in handling insert, update, and delete anomalies, ensuring a better performance of the database.

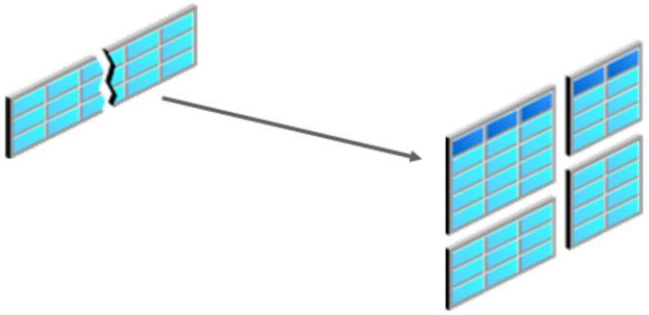

In relational database design, the process of organizing data to minimize redundancy is called normalization. In other words, normalization can be defined as the process of decomposing a table with inconsistencies to produce a smaller, well-structured table. Normalization may involve splitting a table into two or more tables and defining relationships between the tables. The objective is to isolate data so that additions, deletions, and modifications of an attribute or a field can be made in just one table and then propagated throughout the rest of the database by using well-defined relationships. Edgar F. Codd, the inventor of the relational model, introduced the concept of normalization.

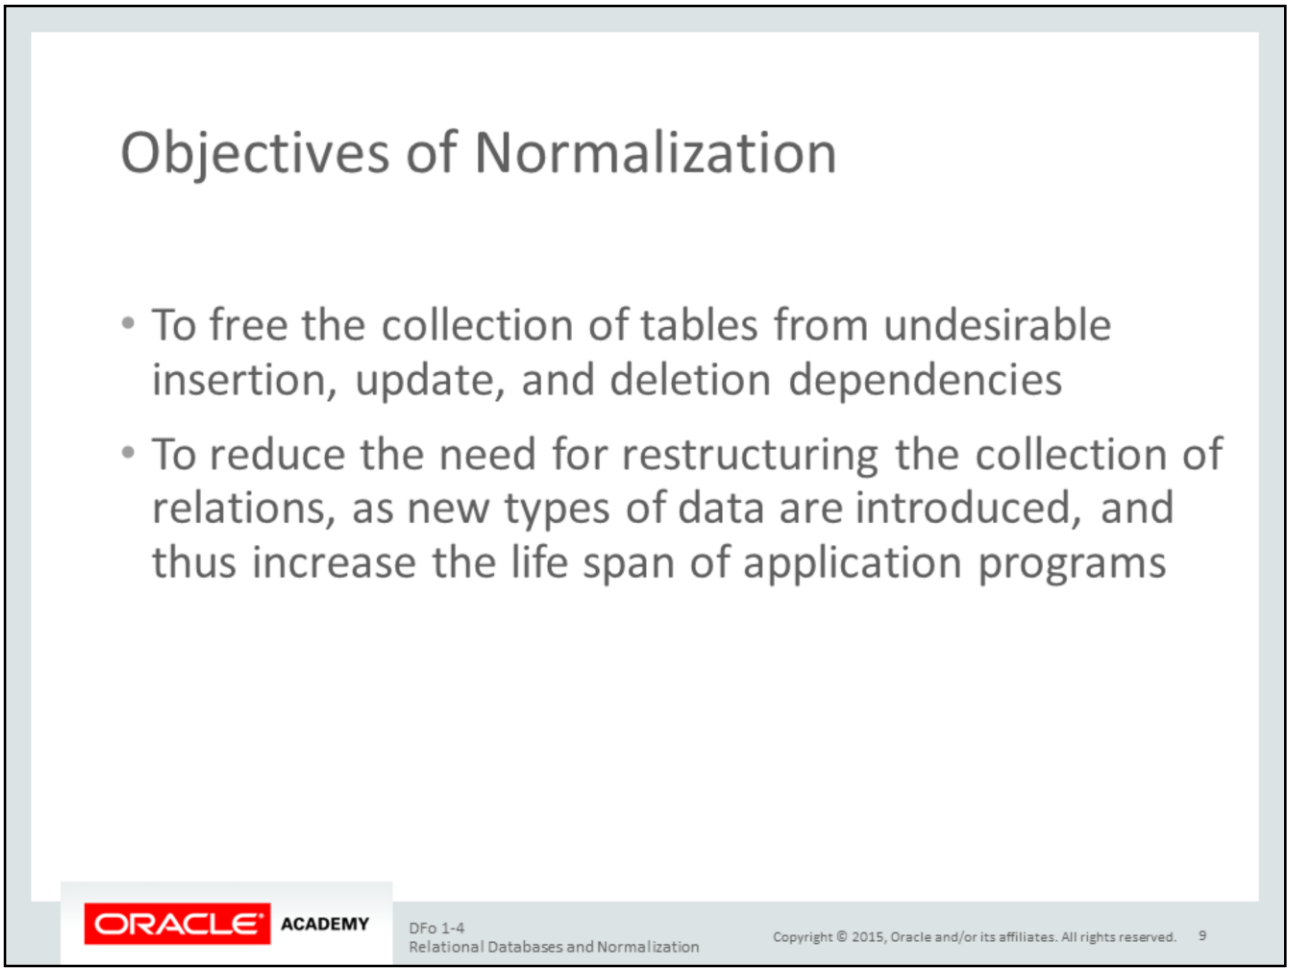

Here are some characteristics of insufficiently normalized tables:

- The same information can be expressed on multiple rows; therefore, updates to the table may result in logical inconsistencies.
- There are circumstances where certain facts cannot be recorded at all.
- Under certain circumstances, deletion of data representing certain facts necessitates deletion of data representing completely different facts.

When a fully normalized database structure is extended to allow it to accommodate new types of data, the preexisting aspects of the database structure can remain largely or entirely unchanged. As a result, applications interacting with the database are minimally affected. Normalized tables are suitable for general-purpose querying; that is, any queries against these tables, including future queries whose details cannot be anticipated, are supported.

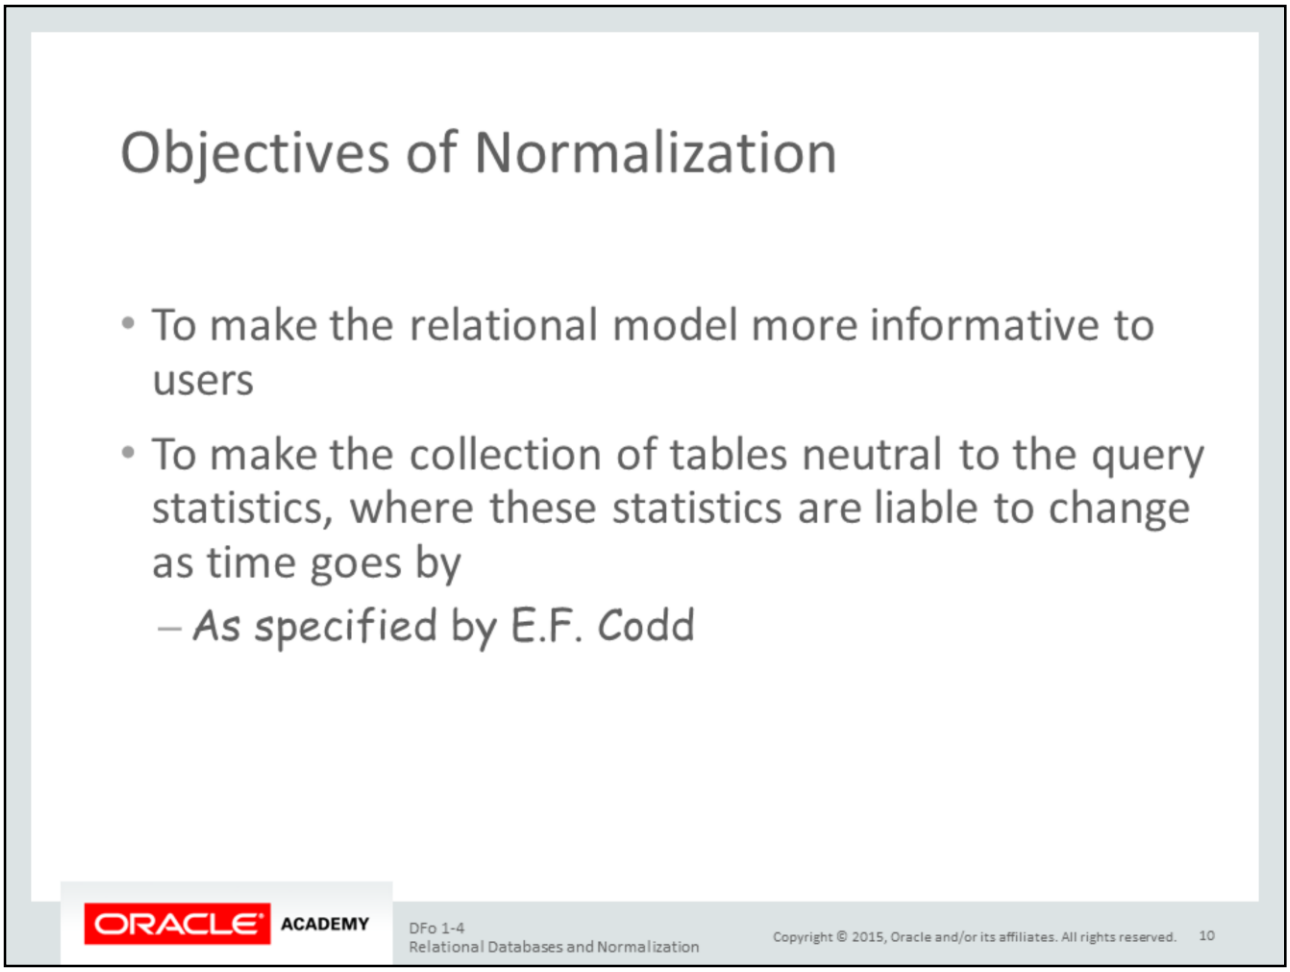

Here are some characteristics of insufficiently normalized tables:

- The same information can be expressed on multiple rows; therefore, updates to the table may result in logical inconsistencies.
- There are circumstances where certain facts cannot be recorded at all.
- Under certain circumstances, deletion of data representing certain facts necessitates deletion of data representing completely different facts.

When a fully normalized database structure is extended to allow it to accommodate new types of data, the preexisting aspects of the database structure can remain largely or entirely unchanged. As a result, applications interacting with the database are minimally affected. Normalized tables are suitable for general-purpose querying; that is, any queries against these tables, including future queries whose details cannot be anticipated, are supported.

Advantages of a Relational Database

- Avoids duplication of data
- Ensures consistency of the data that is stored as records
- Easier to modify data and data format
- Easier to insert and delete data
- Easier to maintain security of data

When you store data in tables, you can easily add, modify, and delete data as well as maintain consistency of the stored information.

Data Integrity

- Data integrity is a very essential function of relational databases.
- Data integrity:
  - Ensures that data is accurate.
  - Ensures that data is consistent.
  - Is achieved through normalization, defined business rules, and validated data.

- There are three main types of data integrity:
  - Entity integrity ensures that each row (record) is a unique instance in a particular table by enforcing the identifier column(s) of a table; for example, Employee ID, Student ID, and Social Security Number (SSN).
  - Domain integrity ensures that valid data is entered for a column by enforcing the data type, the data format, and the range of possible values. For example, the value in the Salary column of the EMPLOYEES table must contain numerical values.
  - Referential integrity preserves the defined relationships between tables when records are inserted, updated/modified, or deleted by ensuring that the key values are consistent across tables. This kind of consistency requires that there are no references to nonexistent values. If a key value changes, all references to that key value change consistently throughout the database.
  - User-defined integrity ensures that the data stored in a database complies with the rules of the business.
- You can also define specific business rules to establish a correct and consistent control of data flow and access in your application.

Section1_Lesson4_Quiz

Quiz - 3 questions

Last Modified: Jul 15, 2016 at 12:49 PM

PROPERTIES

On passing, 'Finish' button:	Goes to Next Slide
On failing, 'Finish' button:	Goes to Next Slide
Allow user to leave quiz:	At any time
User may view slides after quiz:	At any time
Show in menu as:	Multiple items

Edit in Quizmaker Edit Properties

Summary

In this lesson, you should have learned how to:

- Describe the features of a relational database
- Explain the rules of a relational database
- Explain the objectives of normalization
- Describe the types of normalization


ORACLE
ACADEMY